

NBER WORKING PAPER SERIES

STEERING INCENTIVES OF PLATFORMS:
EVIDENCE FROM THE TELECOMMUNICATIONS INDUSTRY

Brian McManus
Aviv Nevo
Zachary Nolan
Jonathan W. Williams

Working Paper 27083
<http://www.nber.org/papers/w27083>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2020

We are grateful for the the North American ISP that provided the data used in this paper. We thank John Asker, Shane Greenstein, Robin Lee, and Jeff Prince for comments, as well as audiences at Indiana-Kelly, Penn State, Yale, FCC, and the NBER's IT and Digitization workshop. Zach Nolan and Jonathan Williams thank Cable Television Laboratories, Inc., and its member organizations, for their support of this work. Brian McManus and Jonathan Williams acknowledge support from the NSF (grant #1824293). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w27083.ack>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by Brian McManus, Aviv Nevo, Zachary Nolan, and Jonathan W. Williams. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Steering Incentives of Platforms: Evidence from the Telecommunications Industry
Brian McManus, Aviv Nevo, Zachary Nolan, and Jonathan W. Williams
NBER Working Paper No. 27083
May 2020
JEL No. L11,L13,L96

ABSTRACT

We study the trade-offs faced by Internet Service Providers (ISPs) that serve as platforms through which consumers access both television and internet services. As online streaming video improves, these providers may respond by attempting to steer consumers away from streaming video toward their own TV services, or by attempting to capture surplus from this improved internet content. We augment the standard mixed bundling model to demonstrate the trade-offs the ISP faces when dealing with streaming video, and we show how these trade-offs change with the pricing options available to the ISP. Next, we use unique household-level panel data and the introduction of usage-based pricing (UBP) in a subset of markets to measure consumers' responses and to evaluate quantitatively the ISP's trade-offs. We find that the introduction of UBP led consumers to upgrade their internet service plans and lower overall internet usage. Our findings suggest that while steering consumers towards TV services is possible, it is likely costly for the ISP and therefore unlikely to be profitable. This is especially true if the ISP can offer rich pricing menus that allow it to capture some of the surplus generated by a better internet service. The results suggest that policies like UBP can increase ISPs' incentive to maintain open access to new internet content.

Brian McManus
Department of Economics CB 3305
University of North Carolina
Chapel Hill, NC 27599
mcmanusb@unc.edu

Zachary Nolan
Department of Economics
419 Chapel Drive, Box 90097
Durham, NC 27708-0097
zach.nolan@duke.edu

Aviv Nevo
Department of Economics
University of Pennsylvania
439 McNeil Building
3718 Locust Walk
Philadelphia, PA 19104
and NBER
anevo@wharton.upenn.edu

Jonathan W. Williams
Department of Economics
107 Gardner Hall, CB 3305
University of North Carolina
Chapel Hill, NC 27599-3305
jonwms@unc.edu

1 Introduction

In a variety of industries, consumers access products or services through a platform, or gatekeeper. When a platform has market power, it faces a trade-off between steering consumers' choices within the platform versus capturing surplus from allowing consumers free choice among the platform's products. On one hand, steering may allow the gatekeeper to direct consumers' choices toward high-margin options, which give it a greater share of the surplus generated on the platform. On the other hand, allowing consumers to have more or better choices may increase the value of the platform, which could allow the gatekeeper to charge more for access. The platform's prices, in addition to directly affecting consumers' welfare, affect the third-party firms that depend on the platform for access to consumers.¹

In this paper, we analyze the steering incentives of an Internet Service Provider (ISP) that sells TV services as well as internet access, and therefore serves as a gatekeeper to online content, some of which is an alternative to the TV service the ISP sells. As streaming services gain popularity and the value of over-the-top-video (OTTV) increases, the ISP faces a trade-off: OTTV directly competes with the TV service but also increases the value of the internet service. The ISP, therefore, needs to choose between trying to steer consumers away from OTTV towards its own video service, or supporting the growth of OTTV in order to improve its internet offering. The ISP's choice of which strategy to follow directly impacts third-party video streaming

¹Examples of industries where gatekeepers are important and might have incentives to steer include the following. In health care markets, consumers enroll in an insurance plan and then access health care providers mostly through the insurer's network. The network's structure and the plan's prices will reflect the insurance company's incentive to steer consumers towards certain providers and reduce moral hazard in health care consumption, while also providing a broad set of care options so that enrollment is attractive for a variety of individuals (Ho and Lee (2019)). Similarly, consumers might reach online shopping sites through a search platform that could have incentives to steer consumers to its own sites or sites that pay a higher fee. Steering activity, however, might turn off some consumers who prefer different products, and therefore reduce total visits to the platform. Indeed, such steering incentives are the basis for the fine imposed on Google in June 2017 by the European Commission. Similarly, in March 2018, Spotify filed a complaint with the European Commission saying that Apple, the gatekeeper of the App Store, "has introduced rules to the App Store that purposely limit choice and stifle innovation at the expense of the user experience – essentially acting as both a player and referee to deliberately disadvantage other app developers." (Spotify CEO public statement March 13, 2019, <https://newsroom.spotify.com/2019-03-13/consumers-and-innovators-win-on-a-level-playing-field/>.)

services, which may lose subscribers, and directly relates to the net neutrality debate. We specifically aim to answer the following questions. What are the forces that drive the trade-off faced by the ISP? How does the trade-off change with the availability of pricing strategies which vary in how the ISP captures value generated by OTTV? And, how are these trade-offs balanced in practice?

To answer these questions we proceed in three steps. First, we utilize a unique household-level panel data set from a North American ISP to document some relevant facts. The data include information on internet usage and TV subscription choices of a large number of households over nearly one year. A unique feature of the data is that we observe which OTTV services a household uses and the volume of traffic generated by these services. Using these data, we show that in our sample, OTTV services (e.g., Netflix, Hulu) account for almost 60% of all data usage on the internet. In addition, we show that households with subscriptions to the ISP's TV service use less data and engage less frequently with OTTV. We document that when a household drops TV service (so called "cord cutting"), it increases its internet usage by nearly one-quarter.

The improvement in OTTV and the cord cutting that follows from it have two implications for the ISP's bottom line. First, a reduction in revenue from TV services, and a loss of positive margins from this service. Second, increased internet traffic that will eventually generate substantial network costs.² This implies that improvements in OTTV quality and access can create an incentive to steer customers towards TV services or even foreclose some OTTV services. These steering and foreclosure incentives are a central concern in discussions of net neutrality and upstream relationships between ISPs and content providers. We show that some of the same impacts can occur through downstream interventions when the ISP steers consumers' choices.

In the second step of the analysis we aim to understand the forces driving the trade-off the ISP faces. To do so, we develop a model of mixed bundling that includes TV and internet subscriptions. We add two elements to the standard bundling framework:

²The increase in network costs is related to OTTV's "unicast" nature, which can have a unique transmission for each consumer

quantity choices by consumers for the services they receive, and a parameter that describes internet subscribers' access to OTTV services. We show that an increase in streaming video access, all else fixed, leads to increased cord-cutting, which creates incentives for the ISP to at least partly foreclose OTTV. However, as we give the ISP the ability to offer internet service tiers that vary by usage allowance, we find that the ISP has less incentive to foreclose and instead it prefers to promote OTTV. The intuition is simple: when the ISP adds internet tiers which associate greater allowances with higher prices, it can extract more revenue from high-demand consumers. Contrary to concerns that usage-related prices or constraints are at odds with net neutrality's goal of robust third-party content creation, we show that ISPs may have an incentive to improve OTTV quality when these pricing tools are available.

In the final step of the analysis, we study the introduction of usage-based pricing (UBP) in a subset of markets. The UBP takes the form of a menu of three-part tariffs, each of which includes an access fee, an internet usage allowance, and an overage charge. The policy's introduction offers an opportunity to measure consumers' responses to steering strategies in this industry and the implications it has for the ISP incentives.

We find the following. First, we find that UBP caused consumers, especially heavy users, to upgrade their internet tiers. This suggests that the ISP can use UBP and a menu of internet plans to separate users according to their internet usage and collect additional revenue from heavy internet users, who value the internet service the most. Second, we find no statistically significant effect on either cord cutting or adding TV service. The relatively muted movement in either direction for TV subscriptions reveals the potential difficulty in influencing consumers along this margin, which implies that ISPs may have broader success with other aspects of steering and surplus capture. Third, we find that UBP caused an 11% reduction in usage, on average, during the treatment period. In both the treated and control markets, usage increased relative to the pre-policy period, however the increase among treated households was smaller. The households that reduced usage by the most, relative to the control markets, had

relatively heavy pre-UBP usage and did not choose to purchase an increased usage allowance. Treated households that upgraded their internet tier, by contrast, used approximately the same amount of data as households not exposed to UBP. Both these empirical findings suggest that the ISP’s introduction of tiers, holding internet content constant, may not lead to pressure for an expanded network and higher ISP investment costs. Finally, our point estimates indicate that UBP led to consumers increasing the video share of their internet usage. This suggests that households favor streaming video over other types of internet content, and therefore reductions in OTTV quality or its foreclosure would not be viewed favorably by consumers.

These outcomes – upgraded internet tiers and reduced overall usage – represent an effective reallocation of OTTV surplus from households to the ISP. This reallocation confirms one of the main channels for incentives in our theoretical model: The ISP’s incentives to encourage new content and expand their networks may be stronger when ISPs capture a greater share of total surplus from OTTV. Overall, these results suggest that while steering consumers towards TV services is possible, it does not seem to be profitable as long as the ISP has the ability to offer different tiers and UBP.

We find that the introduction of UBP did not harm OTTV streaming services and indeed some appear to have benefited. We find that Netflix subscriptions increased when UBP was introduced. Similarly, we find an increase in subscriptions for “linear” OTTV services such as Hulu and SlingTV. We also find that the share of consumption devoted to video increases, although the effect is not statistically significant. Overall, our findings suggest that UBP can simultaneously increase ISP revenue through increased internet usage prices, while also stimulating the OTTV access and usage that are necessary to increase overall surplus from OTTV. Of course, UBP’s impact to reduce overall usage implies that it must negatively affect some internet content, but these effects do not fall on the video sites we highlight in this paper.

Related literature Our analysis relates to previous research on video programming that has focused on traditional TV (Crawford et al. (2017), Crawford and Yurukoglu

(2012), Crawford and Shum (2007), and Crawford and Shum (2015)). Relative to this work, we look at the interaction between TV subscriptions and internet use in general, and OTTV usage specifically.

We follow Nevo et al. (2016), Malone et al. (2016), and Malone et al. (2014), who use high-frequency data to study subscriber behavior on residential broadband networks. These studies cannot separate streaming video activity from other internet usage, and therefore cannot consider ISPs' steering incentives due to OTTV's role in the interaction between internet and television services. Other studies of the demand for and consumers' value from broadband services include Prince and Greenstein (2017), Goetz (2016), Tudon (2018), Goolsbee and Klenow (2006), Dutz et al. (2009), Rosston et al. (2013), Greenstein and McDevitt (2011), Goolsbee and Klenow (2006), Edell and Varaiya (2002), Varian (2002), and Hitte and Tambe (2007).

Relationships between ISPs and internet content providers are an active area for public policy, especially concerning merger approval and net neutrality.³ These policy issues converge with the issue of vertical mergers between ISPs and media companies, which can affect ISPs' profits from various content sources and therefore induce steering activity. The literature on these issues largely began with Wu (2003), who introduced the term "net neutrality" and provides a summary of the issues. Lee and Wu (2009) and Greenstein et al. (2016) discuss and review the literature on the topic. However, most of the existing economic analysis of the topic is theoretical: Economides and Hermalin (2012), Armstrong (2006), Bourreau et al. (2015), Choi et al. (2015), Choi and Kim (2010), Economides and Tag (2012), Gans (2015), Economides and Tag (2016), Reggiani and Valletti (2016), and Sidak (2006). Our empirical analysis on steering incentives complements these theoretical studies by providing insight into relevant elasticities for the debate.

More generally, our model and empirical analysis also contribute to literatures

³The FCC's 2015 Open Internet Order prevented ISPs from discriminating among various online applications. This order limited ISPs' ability to reduce usage of video services from some third-party providers. The FCC voted in 2017 to roll back the order, and future policy in this area continues to be debated.

on firms’ strategic efforts to steer and sort heterogeneous consumers across product menus. Ho and Lee (2019), Liebman (2017), and Raval and Rosenbaum (2017) study how insurers influence patients’ choices across medical providers. Barwick et al. (2017) examine conflicts of interest and steering by residential real-estate brokers. Crawford et al. (2017) consider similar incentives in cable TV markets and estimate the value to cable distributors of including vertically integrated versus non-integrated sports networks in their channel bundles.

The incentive to degrade product quality for discriminatory or steering purposes, as is present in our model, is related to the classic work of Mussa and Rosen (1978), which Crawford and Shum (2007) apply in the context of the telecommunications industry. In the bundling literature, Armstrong (2013) and Gentzkow (2007) study how the consumption of one product in a bundle affects utility from other products, which is similar to the relationship between OTTV and TV that we study. Chu et al. (2011) and Crawford and Yurukoglu (2012) empirically explore how variations on bundling and other pricing strategies can affect firms’ profit and consumer welfare. Nonlinear-pricing strategies similar to those we examine have been studied in broadband markets (Economides and Hermalin (2015), and Lambrecht et al. (2007)), phone service contracts (Miravete (2003), Grubb (2015), and Grubb and Osborne (2015)), and other markets (Hagemann (2017), and McManus (2007)).

2 Data and Industry Background

Internet usage has grown steadily during recent years, largely driven by an increase in streaming video. About 60M U.S. households (46%) used a streaming video service in 2018, up from 44M in 2016 (comScore (2018)). Cisco, a major telecommunications and IT firm, estimates that 81% of North American internet usage was video during 2017, and this share will grow to 85% by 2022 (Cisco (2018)). The emergence and popularity of OTTV services coincides with a trend in consumers dropping their TV service and instead consuming video through the internet, also known as “cord cutting.” Between

2014 and 2017 in the U.S., the number of consumers who cut the cord grew from 3.1M to 14.1M (MarketWatch (2018)). We show that households that drop their TV service increase their internet usage (and similarly households that add TV service decrease their internet usage). We demonstrate the revenue implications of this cord-cutting behavior to motivate the model we present in the next section, which allows us to study ISPs' incentives when faced with this consumer behavior.

2.1 Data

The data we use come from a North American ISP.⁴ Our sample is drawn from a handful of markets during a recent period about a year long, and it is nationally representative in terms of demographics, service offerings, and usage patterns. During our sample period, we observe roughly 350,000 consumers' billing information, subscriptions, and internet usage. To preserve customer anonymity, we do not observe demographic information for individual households.

Like most North American ISPs, this ISP offers internet and TV services via mixed bundling, giving discounts off stand-alone prices when consumers subscribe to both services.⁵ Like other ISPs at the time, the average price difference (across internet tiers) between the bundle and internet-only subscriptions is about \$100. 30% of the ISP customers have internet-only subscriptions, and the remaining 70% subscribe to an internet-TV bundle; no ISP customers in our data subscribe to TV alone. The ISP also offers tiers of internet service differentiated by speed and, as we discuss below, usage allowance in some cases. The typical price difference between adjacent internet tiers is about \$15. In each market we observe, the ISP is the sole provider of high-speed, high-capacity internet service.⁶

⁴Our agreement with the ISP prevents us from identifying the firm or any details that could be used to infer its identity. This includes the specific markets served, the exact dates and details of the implementation of UBP, as well as detailed characteristics of the ISP's product offerings.

⁵The ISP also offers telephone service, which about 40% of its customers subscribe to. We do not use the telephone service information in this paper.

⁶Despite the rapid usage growth within our sample, we see little evidence that congestion affected internet use. Packet loss, which is a quality disruption often caused by congestion, averaged less than 0.01% during peak hours in our sample. See Malone et al. (2016) for a study of the impact of

An important feature of our data is the introduction of UBP in a subset of markets partway through the sample period. We observe data from one treated market and five control markets. The UBP is implemented through a menu of service tiers that consumers can choose from, where each tier is a three-part tariff. Tiers vary in their monthly usage allowance in gigabytes (GBs). Usage up to this allowance is fully covered by the monthly subscription charge, but if a household exceeds its allowance, it is charged a fee which is less than the monthly price difference to move to a higher tier. During UBP, we observe 8% of households across all markets exceed the usage allowance associated with their tier, as it is applied under UBP.

For each household in the sample, we observe download and upload volume for each hour, although for much of the analysis we aggregate usage to the daily level. Households in the sample use a mean (median) of 4.7GB (1.4GB) of internet data per day. Internet usage differs substantially across households. Average usage in the highest-priced (and highest-speed) tier is nearly seven times as in the lowest-priced tier. Within-tier usage dispersion is also substantial; coefficients of variation range from 1.67 to 2.05 across tiers. Across combinations of TV and internet service, internet-only subscribers have mean (median) internet usage 61% (137%) greater than bundle subscribers. Overall, internet usage increased at a 44% annualized rate during our sample period. The growth rates were largest, on average, among consumers who began the sample at lower usage levels, but the increase in total usage volume is driven more by initially high-usage households.

Within a household's total usage, we observe the application (e.g., Netflix) or protocol (e.g., File Transfer Protocol or FTP) generating each byte used by a consumer, but not the specific content (e.g., particular movie title or website).

2.2 OTTV Usage

During 2017, per-household U.S. average daily use of online video was 1 hour and 40 minutes (comScore (2017)). Globally, online video usage grew by 32% per year

congestion on broadband networks.

from 2013-2018 (Zenith (2019)). This growth coincides with the emergence and rapid growth of several prominent firms that offer OTTV services. In 2014 about 40% of US households subscribed to a streaming service such as Netflix (40M subscribers), Amazon Prime (14.5M), or Hulu (7M) (Nielsen (2015)). By 2018, Netflix grew to 60M U.S. subscribers, Amazon to 26M, and Hulu to 27M (eMarketer (2019)).

These trends are reflected in our data. To see this, we divide usage into five major categories: Browsing, Gaming, Music, Video, and Other. Video accounts for 58.5% of all usage, and browsing accounts for another 30.5%. Some browsing activity includes watching videos within web pages, as is possible on Facebook and other popular web sites. The remaining share is divided approximately equally across the other categories. In Table 1, we provide a breakdown of OTTV streaming across specific streaming video applications (HBO, Hulu, Netflix, SlingTV, and YouTube), separately for households who purchase only internet access (top panel) and those who purchase an internet-TV bundle (bottom panel). We classify households as active users of an application at the monthly level, based on whether their application usage is 3GB or greater. Internet-only households are more likely to use each application, and usage intensity of each application is greater as well.

YouTube’s free content is the most widely-accessed video application by households in our sample. Netflix, which offers a variety of original programming along with a library of previously distributed movies and television programs, is the most popular subscription service. HBO Go has a content structure similar to Netflix, while Hulu emphasizes opportunities to stream-on-demand TV shows currently airing on network TV. Other streaming services like SlingTV offer live TV over the internet.

[TABLE 1 HERE]

2.3 Cord Cutting

The emergence and popularity of OTTV services coincides with a trend in cord cutting. In this sub-section we provide a detailed description of cord cutters’ changes in internet

usage.

In Figure 1 Panel (a) we plot the shift in average internet usage by cord-cutting households, and, for comparison, we also plot the average usage of consumers who hold on to their TV subscriptions. The line labeled as “Cord-cutters” is the average per-household usage on each day of the 16-week window centered around the date of subscription change. For comparison, we construct a similar average for all households that did not cut the cord. To construct this average, we calculate the average for each calendar date and then average across the dates to match the windows presented for the cord cutters. We supplement the smooth curves with dashed lines that represent the average levels of internet usage in the weeks before and after cord cutting. Together, these numbers show that the average (eventual) cord cutter’s internet usage begins approximately 31% higher than other subscribers, and increases by 23% after cutting the cord (5.2 GB/day to 6.3 GB/day).⁷ Thus, the individuals who drop TV in our sample tend to be heavier internet users.

[FIGURE 1 HERE]

We explore the nature of cord cutters’ changing internet usage by decomposing the 23% increase in traffic across types of applications. In Figure 1 Panel (b) we provide the change in usage across Browsing, Gaming, Music, Video Streaming, and Other applications. The online activities that increase the most (online video and webpage browsing) provide a fairly similar experience to watching TV. (Some online video delivered through websites, e.g. video on Facebook, appears in our data as browsing rather than video streaming.) Other entertainment applications (Gaming and Music) have increases in usage after cord cutting, as do cloud-based backup services (e.g. Dropbox).

Some changes in Figure 1 Panel (b) are due to increased subscriptions and increased usage of third-party OTTV services. Netflix, Hulu, and SlingTV usage increase by

⁷Our 23% estimate of cord cutting’s impact is likely to be conservative. The data suggest that many consumers “prepare” to drop their TV subscriptions by adding and experimenting with online video services during the weeks before cutting the cord. Cord cutters’ increased usage begins prior to the subscription change date, and consumers who do not alter their subscriptions do not have a corresponding usage change.

36%, 90%, and 538% among households who cut the cord.

[FIGURE 2 HERE]

Consumers who cut the cord sacrifice a portion of their available video entertainment, and in exchange they pay less for the video services they receive. We demonstrate this change in payments by focusing on consumers who begin the sample with a bundle subscription and, at some point, drop TV but remain internet customers. Average payments from these consumers to the ISP fall by about 50%, from about \$130 to roughly \$70. We also calculate estimated total out-of-pocket spending by consumers to the ISP plus all subscription services using 2015 prices. In Figure 2 we plot the densities of consumers' total payments before and after cutting the cord. Average total spending per customer drops by half when the consumer drops a TV subscription. The variation around each distribution's mode is due to differences in consumers' subscription details (i.e., channel packages and internet tiers) and third-party OTTV subscriptions. The revenue reduction is accompanied by an increase in average internet usage, which requires greater data transmission and therefore greater network costs.

The usage, revenue, and cost impacts of cord cutting may be reversed, of course, if the ISP can induce an internet-only subscriber to switch to the bundle. In Appendix Figure A1 we show that households who add TV subscriptions (i.e., "attach the cord") reduce their internet usage by about 12% within two weeks. This reduction is approximately half the magnitude and opposite in sign of cord cutting's impact on usage. We also document in Appendix Figure A1 that OTTV usage falls substantially after a household subscribes to TV.

3 A Simple Model

In this section we introduce a model that allows us to study the trade-offs faced by the ISP in pricing (and offering) internet and TV services. To do so, we augment the standard bundling model in a few ways. First, we allow consumers to make a usage

decision in addition to the subscription decision, so that we capture both intensive and extensive margin choices. Second, we allow consumers with internet subscriptions to access some video content as OTTV. These extensions allow us to explore a range of pricing strategies for the ISP, from simple bundle pricing over subscriptions to more sophisticated nonlinear prices that address consumers’ extensive-margin choices.

We use our model to demonstrate that improvements in OTTV increase the value of the ISP’s internet product, but also increase the ISP’s usage-related costs while reducing demand for the ISP’s TV service, both as a standalone product and in bundles. We use a simple framework and numerical simulations to show that, with bundle pricing only, the ISP steers consumers away from OTTV usage because the firm is hurt by improvements in streaming video. By contrast, when the firm is able to cap internet usage or offer a variety of internet usage tiers, the ISP can benefit from OTTV’s availability. The impact of OTTV on the ISP creates incentives for the firm to expand or restrict third-party firms’ video content, which is central to the net neutrality debate.

3.1 The Setup

Consider a market in which an ISP offers consumers access to two services: non-video internet (1) and video entertainment (2). An individual consumer’s taste for “units” (e.g. hours) of services 1 and 2, relative to the outside option, is given by $v = (v_1, v_2)$. We normalize the consumer population to one and assume that consumers’ tastes are distributed independently and uniformly on $[0, 1] \times [0, 1]$.

To access the services, the ISP offers consumers subscription plans. We begin by assuming that the ISP offers three plans: broadband internet access (i), TV (t), and a bundle (b) that includes both i and t . The firm’s mixed bundling pricing strategy includes prices for the stand-alone plans (p_i and p_t) and a price for the bundle (p_b). A consumer can subscribe to one of the firm’s three plans, $\{i, t, b\}$, or an outside option denoted by 0 that provides utility normalized to zero. To capture the presence of OTTV, we assume that consumers can access some fraction, $\delta \in [0, 1)$, of the

video content (service 2) through an internet-only plan (i). We assume that OTTV is available at no additional expense to the consumer. We initially treat δ as a given, but then explore the ISP's incentives to impact it. The restriction $\delta < 1$ has several possible interpretations, including limited available OTTV content and diminished video quality, which could be due to transmission (e.g. congestion and buffering) or hardware limitations.

Utility

An individual consumer receives utility from consuming q_1 units of service 1 and q_2 units of service 2. The consumer decides on consumption based on his tastes (v) and subscription plan. The quantity choice for video services, q_2 , can include both traditional TV, q_2^t , and OTTV, q_2^i , with $q_2 = q_2^i + q_2^t$. For simplicity, we assume that a consumer has marginal utility equal to one for a service's units up to a satiation level equal to the taste parameter v , and then marginal utility is zero for any greater quantity.⁸ For example, a consumer with taste v_2 and a TV-only subscription consumes $q_2 = q_2^t = v_2$ units of video entertainment through his TV and receives surplus of v_2 from this activity. When the consumer uses streaming video, his marginal utility from video hours remains equal to one up to δv_2 , where it falls to zero. This can be viewed as a scenario where a consumer enjoys v_2 distinct shows available on TV, but only the fraction δ of the shows are available through OTTV. To simplify the consumption choices of bundle subscribers, we assume that TV subscribers receive all video entertainment through TV, with $q_2^t = v_2$ and $q_2^i = 0$.⁹

Putting this all together, subscribers in internet-only plans receive utility of $U_i = v_1 + \delta v_2 - p_i$, where the first and second terms capture utility (and quantities) from consuming non-video internet and OTTV applications, respectively. A subscription to the TV service, t , results in video consumption of $q_2^t = v_2$, zero non-video internet

⁸This structure may arise if conventional TV carries v_2 shows that each yield utility of one, while all other shows yield a utility of zero.

⁹This assumption represents the best-case/least-cost outcome for the ISP and strengthens incentives to steer consumers to the bundle, but the assumption does not qualitatively change any of the model's predictions.

usage given the lack of access, and net utility equal to $U_t = v_2 - p_t$. Bundle subscribers consume quantities of video and non-video internet up to their satiation levels and receive utility equal to $U_b = v_1 + v_2 - p_b$. Finally, if the consumer selects the outside option, 0, quantities are zero for both services and utility is zero.

The ISP Profits and Costs

The ISP's profit from serving the market depends on the number of subscribers to each plan, the quantities they select, and the costs of providing these services. For simplicity, we assume that the ISP's costs depend only on the quantity of internet content it transmits; all other costs are zero. This means that we abstract-away from costs such as affiliate or retransmission fees paid to content firms for individual TV subscribers, or the fixed costs of operating the ISP's network.¹⁰

The ISP's internet-related costs increase with the bytes transmitted through its network and therefore are increasing in total usage, which is a function of the number of users and usage (and usage type) by each user. These costs capture the ISP's need to maintain a network that must grow with total usage to avoid congestion, content interruptions, or slowdowns during times of high total usage. We assume that the firm's transmission costs are given by a parameter, γ , times the gigabytes, g , required to transmit the quantities q_1 and q_2^i . We assume that one unit of non-video internet content requires the transmission of one gigabyte, so $q_1 = g$. We further assume that consumption of video content q_2^i requires βg gigabytes, with $\beta > 1$. This captures the greater bandwidth and immediacy needs to transmit video. Consumers' preferences and their choices, given the firm's prices, imply a mapping from tastes to quantities of internet usage, so we write a consumer's internet usage as a function of v . Without restrictions on usage, an internet-only subscriber uses $g_i(v) = v_1 + \beta\delta v_2$ internet units, while a bundle subscriber uses $g_b(v) = v_1$.

Combining consumers' choices with the firm's cost structure, the ISP's profit function is:

¹⁰Our model's qualitative results would not change if we were to include positive per-subscriber costs for TV service. See additional discussion below.

$$\pi = \int_{v \in S_t} p_t dv + \int_{v \in S_i} [p_i - \gamma g_i(v)] dv + \int_{v \in S_b} [p_b - \gamma g_b(v)] dv.$$

The terms S_t , S_i , and S_b are the sets of consumers who choose internet, TV, and the bundle, respectively.

3.2 Consumer Choice

We now turn to the choices consumers make in this setup. In Panel (a) of Figure 3 we present choices different consumers make for a fixed set of prices. When video can be consumed only through TV (i.e., $\delta = 0$), consumers in the areas labeled ‘0’ and ‘I’ select the outside option, and those in areas ‘b’ and ‘II’ select the bundle. Consumers in areas i and t select the stand-alone internet and TV subscription plans. The split is intuitive: consumers with low valuations of both products choose the outside option. Consumers with high valuations for both products choose the bundle, and consumers with high valuation for one service and not the other choose the plan with just one service. The exact boundaries between the different choices depend on the prices of the various options.

[FIGURE 3 HERE]

We next consider the effect of OTTV becoming more attractive (i.e., the effect of an increase in δ) on choices. In Panel (a) we show the effect of δ increasing to 0.7, holding prices fixed at the optimal level for $\delta = 0$ and $\gamma = 0.2$. Two types of consumers change their choices. First, some consumers (in area I) who did not purchase, despite moderately high valuation for internet or video, will now subscribe to i because it became more attractive. These new consumers can increase the ISP’s profit and are one reason the ISP has an incentive to encourage OTTV. Second, some consumers (in area II) decide to “cut the cord.” These consumers initially choose a bundle, but have relatively low video entertainment tastes among bundle subscribers. As δ increases, they prefer stand-alone internet service because they can consume video over-the-top using the internet service. The cord cutting by these consumers diminishes the ISP’s

revenue, since the bundle price is higher than the internet-only price, and increases the ISP's costs, since supplying video over the internet is costly to the ISP.¹¹

These two groups of consumers have opposite effects on the ISP profits, and represent the basic trade-off the ISP faces: stronger OTTV makes the internet service more valuable but also competes with the TV service. In principle, which effect dominates depends on, among other things, the distribution of tastes. In Panel (b) of Figure 3 we show the effect when tastes are uniformly distributed. We plot market shares for each plan at different values of δ , holding prices fixed. The share of internet-only plans, i , grows monotonically in δ , and this largely comes from a reduction in the bundle market share, which is completely eliminated around $\delta = 0.85$. When OTTV is sufficiently beneficial, some consumers even convert their t subscriptions to i , as i offers services very similar to the original bundle.

Under our distributional and cost assumptions, ISP profits fall when OTTV improves, holding prices fixed. Many consumers “cut the cord” and trade a more expensive subscription for one with a lower price. This reduces the ISP revenue and increases its costs from video usage previously delivered at no marginal cost. The modest increase in total subscriptions is not enough to make up for the decrease in profit from reduced bundle subscriptions.

3.3 ISP Pricing Strategies

In the previous section we showed that an improvement in OTTV can lead to “cord cutting” and a loss in ISP profits, holding prices fixed. Of course, in equilibrium the ISP can and will respond. In this section we study some possible responses, including changing the relative prices of products, limits on internet usage, and tiered pricing of internet usage. Our discussion in this section takes δ as fixed, but in the following section we consider the ISP's incentive to increase or decrease δ , which it could do by promoting or restricting access to OTTV.

¹¹The ISP will lose profit on these consumers even with a positive cost per TV subscriber, as long there is a positive margin on the TV part of the bundle (specifically, the difference between the bundle price and the internet-only price is greater than the per-TV subscriber cost).

Subscription Prices and Bundle Discounts

In Figure 4 we illustrate the effects of allowing the ISP to set profit-maximizing bundling prices for different values of δ . In Panel (a) we display the prices, while in Panel (b) we present the market shares of t , i , and b given these prices. By comparing Figures 3 and 4, we can see clearly the steering effect prices can have. Greater values of δ drive the ISP to set greater values of p_i and greater bundle discounts, i.e. the difference between p_b and $p_i + p_t$. By increasing the price of the internet-only product, the ISP can push back and reverse the effect of an increase in δ on cord cutting. As we show in Panel (b) of Figure 4, with optimal mixed bundling prices the share of the internet-only plan actually decreases. The steering effects are especially strong for $\delta > 0.8$, when the optimal pricing sets $p_b = p_i$, and therefore the share of internet-only subscriptions goes to zero.¹² With the pricing instruments available in this case, the cost impact of δ is greater than ISP's opportunity to profit from the increased value of internet-only subscription plans, so it effectively shuts down this plan.

[FIGURE 4 HERE]

In the particular setting we evaluate above, we find that optimizing prices mitigates the decline in profits but does not eliminate it completely. We discuss below the impact of this pricing strategy and others on ISP profit.

Usage Limits

As we showed in the discussion above, firms may find it optimal to set prices such that the internet-only product is basically eliminated. However, if firms are permitted to use more sophisticated strategies, their incentives to steer consumers into the bundle may be reduced. One possible strategy is to charge consumers an access or subscription fee in return for a usage allowance which puts a strict cap on internet usage. This is a simplified version of the three-part tariff we see in our data, where the consumer pays an overage charge when he exceeds the allowance. With our assumption of marginal

¹²Under our assumption that bundle subscribers receive all of their video entertainment via the TV subscription, the elimination of i also eliminates OTTV usage. In a more realistic model, bundle subscribers would consume video through both TV and the internet, and OTTV consumption would not be eliminated.

utility equal to either one or zero, depending on the quantity consumed, an optimally-set overage charge is equal to one, and ISP's only interesting strategic choice is the usage allowance.

Let κ denote the usage allowance in GBs. The effect of the allowance is straightforward: as long as it is set low enough, the allowance will truncate the usage of any consumer with $v_1 > \kappa$, regardless of whether they purchase a bundle or the internet-only plan. This allows the ISP to keep cost down even as δ increases; namely, the ISP can limit how much consumers can increase their use of streaming video when δ improves.

In Figure 5 we illustrate the consequences of allowing the ISP to set κ . In Panel (a) we illustrate optimal prices and allowances for a range of δ values, and in Panel (b) we present the resulting market shares. The allowance has a significant impact on the ISP's choices of subscription plan prices and therefore market shares. The internet and bundle prices both increase gradually with δ , and internet-only subscriptions eventually overtake the bundle as the most popular option. Overall, the ISP serves a greater share of the full consumer population as δ rises. The usage cap, which is just above p_i for low values of δ , increases more quickly with δ than p_i , which effectively loosens the ISP's restrictions on internet usage as streaming video improves. (Some i subscribers with large v_1 and small v_2 use no OTTV, while others consume both internet services and streaming video.) Thus, introducing the allowance can permit the ISP to balance providing access to more consumers while also steering customers with greater video demand toward the bundle, where they can receive video at lower cost to the ISP. A usage cap can actually lead the ISP to permit more internet and OTTV usage than it would if usage caps were prohibited, because the cap allows the ISP to influence usage more selectively.

[FIGURE 5 HERE]

Usage Tiers

In addition to using base prices (p_t, p_i, p_b) and allowances (κ) to steer consumers

among packages, ISPs also often separate their internet services into tiers that vary in allowances, overage prices, or connectivity speeds. We consider the effect of the ISP introducing a “low-usage” internet tier with a usage allowance κ and a “high-usage” tier with no usage limit. Consumers can select one of these internet tiers alone or as part of a bundle with conventional TV. For simplicity, and consistent with industry practice, we assume that the difference (i.e. discount) between the bundle price and the sum of TV and internet prices is the same regardless of which tier a consumer chooses.

When the ISP introduces tiers and allowances to a setting which had neither, several types of consumers are affected in distinct ways. First, some consumers who had internet-only subscriptions will stay with this type of plan, with consumers separating between capped and uncapped tiers based on their valuations of internet content (including OTTV). Second, some internet-only consumers with strong video tastes will switch to a bundle plan, with the internet tier depending on their value of v_1 . Finally among initial i subscribers, some will cancel their subscriptions completely because their relatively strong taste for video entertainment and weak taste for non-video internet means that capped internet and bundle subscriptions are worth less than the outside option. Turning to consumers who previously choose a bundle, some with strong internet tastes will opt for the unlimited usage option so that they can consume more content. In the Appendix we provide additional discussion and graphical illustration of these effects.

[FIGURE 6 HERE]

Next, we consider how the ISP would set prices and tiers for different values of δ . In Panel (a) of Figure 6, we provide the ISP’s optimal prices and cap values, and in Panel (b) of Figure 6, we show the resulting market shares. In this setting, the ISP has several levers for steering consumers’ choices, extracting rents, and reducing costs, and as a consequence the ISP’s strategy varies qualitatively with δ . For δ values less than about 0.5, the ISP allows prices to increase with δ so that it captures some of the additional value available to consumers through OTTV. (The high-usage bundle’s

share falls in δ because video usage opportunities remain constant for TV subscribers.) For relatively high values of δ , the ISP's strategy changes, and it raises the high-usage internet price substantially to shrink this plan's market share. When δ is large, the cost of offering unlimited internet usage is too great relative to the premium the ISP can charge consumers for the plan. As the ISP reduces the high-usage internet tier, it raises the low-usage tier's allowance, which accommodates more OTTV usage within this tier. The ISP benefits, on the whole, from steering consumers into the low-usage internet tier where it can recover value from OTTV with an increased price, while also limiting the usage of inframarginal consumers whose OTTV usage choices, if unchecked, would be costly for the ISP.

3.4 ISP Incentives to Promote or Restrict OTTV

We now explore the ISP's direct incentives to either promote or restrict the quality of OTTV by altering δ . For example, the ISP Comcast's Xfinity X1 platform includes one-touch access to Netflix through its remote control. In our analysis of the three pricing strategies discussed above—bundle pricing, usage limits, and usage tiers—the ISP's approach to OTTV varies substantially. With bundle pricing, OTTV usage first grows with δ and then falls as the ISP effectively eliminates its internet-only subscriptions. With usage limits and tiers, by contrast, the ISP allows OTTV usage to grow with δ within significant internet-only market shares. This divergence in OTTV usage occurs because the more sophisticated pricing strategies allow the ISP to capture some of the value generated by OTTV while limiting its effects on the ISP's costs and TV demand. The same logic holds in the results below.

To demonstrate these effects, we plot in Figure 7 the ISP profits for different pricing strategies and different values of δ . These profits correspond to the prices and market shares presented in Figures 4, 5 and 6. As expected, Figure 7 shows that the firm's profit is weakly greater for any δ when it has more instruments in designing a pricing plan. More importantly, the strategies' profit levels vary in whether they increase or

decrease in δ . The variation in the slope across strategies comes from differences in the ISP's ability to extract rents versus avoid additional costs when OTTV's quality improves.

[FIGURE 7 HERE]

For the simplest pricing strategy (bundle pricing), Figure 7 shows that the ISP's profits decrease mildly in δ . If the firm can set a usage cap, by contrast, its profit increases in δ . With the cap, the ISP receives the revenue benefits of increased prices (due to increased value to consumers) from consumers who would use a low or moderate amount of OTTV even without a cap, plus it can limit the OTTV-related costs of internet-only consumers who would use larger quantities of streaming video without a cap. Finally, when the ISP can combine a cap and tiers, its profit may not be monotone in δ . In the example illustrated in Figure 7, the change in profit's slope coincides with the ISP's qualitative change in pricing strategy. When δ is zero, the ISP benefits strongly from the opportunity to segment consumers by their taste for non-video internet, but once δ increases there are additional costs from uncapped internet-only subscriptions. These costs reduce ISP profit until the firm (largely) turns to steering its internet-only consumers into capped subscriptions, where the ISP's price and profit rise as OTTV improves.

The differences in profit slopes and levels across δ values come from the firm's opportunity to balance a complex set of trade-offs regarding consumer choices. Our model shows that, even in a simple framework, an ISP's policies toward OTTV will depend on the current quality of streaming video and what pricing tools ISPs may use. ISPs restricted from setting overage charges or creating tiered service might often benefit from stifling OTTV access, while ISPs with more pricing tools may encourage OTTV usage. More generally, our model demonstrates that prices and plans offered to consumers are a key component of the net neutrality debate; restrictions and promotion of third-party internet content do not just happen at the level of ISP-to-firm negotiations upstream.

4 Lessons from the Introduction of UBP

The theoretical model in the previous section highlights the ISP’s trade-offs when faced with new streaming services and improvement in OTTV quality. If the ISP uses prices to steer consumers toward its own TV services, it gives up the opportunity to capture surplus from the improvement in OTTV. On the other hand, if the ISP allows consumers to freely use OTTV, demand for internet subscriptions will increase but the ISP’s costs increase as well. Our theoretical model shows that the ISP’s overall profits from OTTV, and therefore whether it prefers to encourage or discourage streaming video, depend on consumers’ preferences and the ISP’s ability to recover some of the value generated by OTTV.

In this section, we utilize a change in internet usage pricing to generate evidence that informs the theoretical discussion. As we described in Section 2, we observe data from one market where the ISP introduced UBP—a three-part tariff on internet usage. We use this introduction of UBP to learn about the willingness of consumers to substitute along various dimensions. This is informative about the ISP’s ability to steer consumers’ choices and capture surplus from internet usage, and therefore the ISP’s incentives to restrict or promote OTTV.

After introducing the UBP policy in section 4.1, we answer the following questions in section 4.2: Do consumers upgrade their internet tiers? Do they add or drop TV service? Do they reduce internet usage in general and OTTV specifically? Do these responses vary between heavy and light users? In section 4.3 we consider the implications of the answers to these questions for ISPs and third-party OTTV providers.

4.1 The Introduction of UBP

We observe the introduction of UBP in a “treated” market plus contemporaneous data from “control” markets where the introduction did not happen during the sample period. The ISP’s introduction of UBP came in two phases. The ISP announced that UBP would be implemented starting on a specified date, and provided households

with information about how their monthly usage compared to the data allowance within their current internet tier. During this phase, which we call the “announcement period,” households were not billed if their usage exceeded their tier’s allowance. In the next phase, which we call the “treatment period,” the ISP assessed overage charges on households that exceeded their allowances. We observe several months of activity (a “pre-policy period”), immediately prior to the announcement and treatment periods, which each last a few months.

During the pre-policy period the treated and control markets were subject to the same pricing policies by the ISP. Other than the introduction of UBP, the ISP’s internet tiers were identical across markets in our sample. To our knowledge, the treated market we observe was chosen “randomly” by the ISP as a trial setting, before UBP was introduced more broadly. Across all markets we observe, competitors’ subscription offerings did not change meaningfully during the sample period, including in response to the UBP policy’s introduction. Satellite TV was available in each market, as was internet service via DSL lines. During our sample period in the markets we study, the ISP offered internet service on a substantially higher-speed and higher-capacity network than the alternatives.

To evaluate the comparability of treated and control markets for the empirical analysis described below, we examine pre-policy trends for a collection of variables that are central to the analysis that follows. Specifically, we focus on TV subscription and internet tier transitions, overall internet usage, and the shares of usage devoted to video and prominent OTTV services, and frequency of engagement with OTTV services. For these variables, we test whether there are differences in propensity to subscribe to a service, or different pre-policy trends in the treated market compared to the control markets. The results are presented in the appendix. We generally find no meaningful economic differences between the sets of markets.

4.2 Households' Responses to the Introduction of UBP

In this subsection we analyze households' responses to the introduction of UBP along various dimensions of choice.

4.2.1 The Impact on Choice of ISP Services

We start our analysis by studying UBP's impact on the ISP services that households choose. Specifically, we focus on the choice of internet tier and whether to purchase TV access. We compare the changes in the choices of households in the treated market, where UBP was introduced, to changes in the control markets.

We utilize a Probit model to capture the probability that a household h changes a subscription decision, denoted by $y_h = 1$. We study various decisions, and in all cases we ask whether the household changed the service between the end of the pre-policy period and the end of the sample period. In other words, an observation is a household and the dependent variable is equal to one if a household made the decision studied, regardless of whether this occurred during the announcement or treatment periods. Formally, the model is given by:

$$Pr(y_h = 1) = \Phi(\beta_0 + Treatment_h \times (\beta_1 + \sum_N ShareN_h \times \beta_{2,N}) + \theta X_h), \quad (1)$$

where Φ is the CDF of the standard normal distribution. $Treatment_h$ is an indicator variable equal to 1 if household h resides in a market where the ISP introduced internet usage prices. $ShareN_h$ is a dummy variable equal to 1 if pre-policy household h usage was in range N , where N is below 50%, between 50% and 100%, or above 100% of household h internet tier's treatment-period allowance.¹³ The $ShareN_h$ variables capture the households' likelihood of receiving an overage charge due to usage beyond

¹³Across all households, 80.5% have a usage share less than 50% of the allowance, 15.1% have usage between 50% and 100%, and 4.3% have pre-policy usage greater than 100% of the eventual allowance. Among households that subscribe to Netflix or a linear OTTV service during the pre-policy period, the proportions of households in the same usage share categories are 71.4%, 22.0%, and 6.6%, respectively.

the tier allowance. Finally, all specifications include controls in X_j for households' pre-policy internet usage levels divided into deciles, their pre-policy internet usage growth interacted with usage deciles, and their initial ISP and OTTV subscriptions.

In Table 2 we report the marginal effects from estimates of equation (1). In columns (1) and (2) we ask whether a household upgraded the internet service tier after the UBP policy was introduced. The estimation sample includes all observed households except those already in the top tier at the end of the pre-policy period. The upgrading decision directly informs whether an ISP can induce internet users, in general, and heavy users in particular, to pay tier-upgrading fees, in which case the ISP may have less incentive to reduce access to OTTV. In the raw data we see that (between the end of the pre-policy period and the end of the sample) 7.9% of households in the treated market upgraded their internet tier compared to 4.0% in the control markets ($p < 0.01$). In column (1) we show that this difference persists once we control for observable differences across households: households in the treated market are 3.0 percentage points more likely to upgrade. This difference implies a 0.3% increase in ISP revenue from tier upgrades alone.

In column (2) we ask whether there is heterogeneity in the effect. The theoretical prediction, in Section 3, suggests that usage prices induce more tier-upgrading activity by households who are heavier users. The estimates are consistent with this prediction; households' tier upgrading increases in their average pre-policy usage. Households who use less than 50% of their allowance in the pre-policy period are about 1.6 percentage points more likely to upgrade their tier, while households with usage shares between 50% and 100% upgrade 5.3 percentage points more frequently, and households with usage shares above 100% upgrade 11.2 percentage points more frequently.¹⁴

[TABLE 2 HERE]

In columns (3) and (4) we study the decision to downgrade the service tier. We study this choice among all households except those in the lowest tier at the end of the

¹⁴In additional unreported analysis, we investigate whether upgrade probabilities vary significantly for Netflix or linear OTTV subscribers. They generally are not.

pre-policy period. In the data, 2.9% of households in the treated market downgraded their service, compared with 2.6% in the control markets. This result is unchanged once we control for cross-market observable differences, in column (3). We find that UBP generates a 0.2 percentage point increase in downgrades, which implies about 0.02% less subscription revenue for the ISP. In column (4) we show that the households most likely to downgrade used relatively low shares of their tier allowances. For these households, the introduction of UBP may have increased awareness that a lower internet tier would accommodate their likely amount of usage.

Next, we turn to the impact of internet usage prices on households' TV subscription choices. In columns (5) and (6) we analyze whether the introduction of UBP was successful in reversing cord cutting by asking if it steered internet-only households into the TV-internet bundle. The relevant sample is the households who did not have a TV subscription at the end of the pre-policy period. We see in the data that 4.1% of households in the treated market add TV, compared to 3.7% of households in the control markets ($p = 0.10$). Once we control for differences across household, in column (5), we estimate that treated households are 0.3 percentage points more likely to add TV subscriptions; this estimate is not statistically significant. This modest change to TV subscriptions, which increases the treated market's revenue by about 0.05%, may be due to the relatively high price of adding TV (about \$100) compared to the household monitoring its internet usage and paying for a tier upgrade or occasional overage fee. As we show in column (6), there are no values of $ShareN_h$ for which there is a significant difference between the treated and control markets.

We generally find similar results for cord cutting, focusing on households that ended the pre-policy period with a TV subscription. In both the treated and control markets, 2.6% of households cut the cord. As we show in column (7), the difference remains close to zero when we control for observable differences across households. We report similar results in column (8), which allows for potentially heterogeneous treatment effects across internet usage intensity. Treated households with usage shares between 50% and 100% are slightly more likely to drop TV, but we find a similar shift away

from cord-cutting by treated households in the highest usage group.

4.2.2 The Impact on Internet Usage

Next we study how the introduction of UBP impacted internet usage. To do so, we estimate the following model:

$$\begin{aligned} \log(q_{h,t}) = & \beta_0 + Treatment_{h,t} \times (\beta_1 + \sum_N ShareN_h \times \beta_{2,N}) + \theta X_{ht} \\ & + \tau Trend_{h,t} + \eta_h + \epsilon_{h,t}, \end{aligned} \tag{2}$$

where the dependent variable is the logarithm of household h internet usage in month t . $Treatment_{h,t}$ is a dummy variables equal to 1 if h was in the treated market and t was during the treatment period. As in the analysis in the previous sub-section, we interact $Treatment_{h,t}$ with $ShareN$ indicators. The vector $X_{h,t}$ includes the dummy variable $Announce_{h,t}$, which is defined similarly to $Treatment_{h,t}$ but for the announcement period, plus the interaction of $Announce_{h,t}$ with $ShareN$ indicators. We also allow for household fixed effects (η_h) and cubic time trends that vary with deciles of pre-policy internet usage.

We report the estimation results in Table 3. In columns (1) and (2), we estimate the model using data from all households. In column (1), we find that UBP caused an 11% reduction in usage during the treatment period. In both the treated market and the control markets usage increased relative to the pre-policy period, however the increase in the control markets was smaller. In column (2) we allow the effect to vary with the pre-policy usage of the household. We find that the reduction in usage is significant for both light and heavy users, but the reduction is greater for heavy users within a tier. Households in the two lowest usage categories reduced usage by about 10%, while households with usage shares above 100% made usage reductions nearly twice as large.

[TABLE 3 HERE]

To explore UBP’s impact further, we present separately, in columns (3)-(6), the

effects for households who upgraded their tier and those who did not. Incentives to reduce usage are different between the two groups. We find that households who remained in their original tiers reduced usage by 12.1%, while those who upgraded reduced usage by only 4.4%, which is not statistically significant. This difference in usage responses is consistent with our theoretical model, and it highlights an important channel for usage prices to increase ISP profits. Households have the option to reduce usage (thereby reducing costs) while keeping subscription payments fixed, or they can upgrade and continue their previous usage habits while allowing the ISP to capture a greater share of the surplus. The differences are even larger when we allow the treatment’s impact to vary with the pre-policy usage.¹⁵ We find that households who did not upgrade have larger reductions in internet usage for each value in $ShareN$, including a 33% reduction in usage by households in the highest category.

In Table 4 we report estimates that explore the source of the reduction in usage. We do so by estimating a version of equation (2) where we regress the monthly share of household’s h internet usage that is video in month t on the $Treatment_{h,t}$ and $Treatment_{h,t} \times ShareN$ interactions. We include the same collection of control variables as in equation (2), including household-level fixed effects. This sample includes fewer observations than Table 3 because we do not have complete data on linear OTTV usage from all control markets during the treatment period.

[TABLE 4 HERE]

Overall, once we cluster the standard errors at the market level, the results are statistically insignificant. In what follows we discuss the point estimates. In column (1), the point estimate suggests that for the average user the effect of UBP is an increase in video share by 1.2 percentage points on a base of roughly 38%.¹⁶ This suggests that households place greater priority on streaming video than other types of

¹⁵Interestingly, the effect among households who upgrade is not uniform. The heaviest users upgrade and usage grows more slowly than untreated households. The users who previously used between 50% and 100% of their allowance upgrade and their usage increases more rapidly than similar control households.

¹⁶This figure represents the video usage share of the average user, which is lower than the average share of video in usage since heavier users have higher usage and a much higher video share.

internet content. This may reflect how difficult it is to move households from internet-only subscriptions into the bundle. In column (2), when we separate the effect of the treatment by pre-policy usage share, the point estimates suggest that higher-share households increase their video share by larger amounts. Higher-share households generally use more video, but the treatment’s impact is greater in percentage terms for these households. In column (3), we limit the sample to household-month combinations during which we infer that the household has a Netflix or linear OTTV subscription,¹⁷ and we find results similar to those presented in column (2). In columns (4)-(7) we examine the share of all usage that is devoted to Netflix or linear OTTV. Columns (4) and (6) include all households in the analysis, while columns (5) and (7) condition on a household having a Netflix or linear OTTV subscription, respectively. In all cases, the point estimates suggest that household internet usage becomes more focused on these large OTTV services under UBP.

4.2.3 The Impact on Subscriptions to Third-Party OTTV Providers

To further investigate the effect of UBP on video usage, we examine the impact on third-party OTTV firms through households’ subscriptions. We focus on Netflix and linear OTTV (Sling or Hulu) services. To measure the impact of UBP, we ask whether a household subscribed to a particular service during the treatment period. As in the earlier analysis of ISP subscriptions, we estimate this effect using a Probit model with one observation per household. Specifically, we estimate a version of equation (1), where the dependent variable is whether the household subscribes to an OTTV service during one or more treatment period months. As in the analysis in Table 4, we limit the sample to the markets in which we have complete data for all major OTTV services. We present the results in Table 5.

We find that Netflix subscriptions increase by 3.3 percentage points in the treated market relative to the control markets. The response varies somewhat by pre-policy

¹⁷We infer Netflix and LinOTTV subscriptions from the observed usage data. Specifically, if the household uses more than 3GB of Netflix or LinOTTV in a month, then we say that it is a subscriber in that month.

usage. In column (2) we focus on households that did not have a Netflix subscription during the any month of the pre-policy period (and ask if they added Netflix), and in column (3) we focus on households that subscribed to Netflix prior to UBP (and ask if they kept the Netflix subscription). Households with low usage shares were more likely to add (column (2)) or retain (column (3)) subscriptions in the treated market, while higher-usage households in the treated market were less likely to add a Netflix subscription under UBP. This aligns with the price incentives of UBP; low-share households are the least likely to incur usage-related overage charges due to new subscriptions.

[TABLE 5 HERE]

We report results focusing on linear OTTV subscriptions, instead of Netflix, in Table 5 columns (4)-(6). As in the case of Netflix subscriptions, linear OTTV subscriptions increased in the treated market. The overall effect of UBP was a 0.2 percentage point increase in subscription frequency; this change is on a base of about 8% during the pre-policy period. The greater frequency of linear OTTV subscriptions under UBP is largely due to households who were pre-policy subscribers and kept their linear OTTV service during the treatment period. As in the case of Netflix, we find the linear OTTV subscriptions are more common among households with usage shares below 100% of pre-policy tier allowances; the coefficient estimates, however, are generally not significantly different from zero.

The increases in Netflix and linear OTTV subscriptions, combined with the increase in streaming video usage documented in Table 4, suggests that the ISP's UBP policy did not come at the expense of OTTV demand. In fact, UBP appears to have prompted treated customers to focus on these high-profile third-party OTTV firms. The UBP policy, therefore, can simultaneously increase revenue through increased internet usage prices while also stimulating the OTTV access and usage that are necessary to increase overall surplus from OTTV. Of course, UBP's impact to reduce overall usage implies that it must negatively affect some internet content, but these effects do not fall on the video sites we highlight in this paper.

4.3 Implications for ISPs' incentives for OTTV

Broadly speaking, an ISP has two possible responses to new streaming services or more generally an improvement in the quality of OTTV. One possible response is to steer consumers away from these services and toward the ISP's own TV products. Alternatively, the ISP may decide to promote OTTV and enhance the quality of the internet offering. As we show in Section 3, this strategy is particularly profitable if the ISP can extract surplus from consumers who place higher value on OTTV using some form of price discrimination or nonlinear pricing. The former strategy is associated with reducing OTTV usage and possibly access, while the latter benefits from allowing third-party video content to grow. ISPs' choices between these strategies will depend on the pricing tools available and the difficulty or ease of moving consumers across products.

In Section 3, we provided simulation results for particular parameter values and specific distributions of tastes that suggest that while the first option (of steering consumers towards TV services) is possible, it does not seem to be profitable if the ISP has sufficient pricing tools such as the ability to offer different tiers and UBP. The empirical results in the previous subsection support the same conclusion for the following reasons.

First, the results in Table 2 columns (1) - (4) show that the ISP can use UBP and a menu of internet plans to induce consumers to change their internet usage allowance. This allows the ISP to separate users according to their internet usage and collect additional revenue from heavy users, who value the internet service the most. Collecting more revenue makes the option of promoting OTTV more profitable because the ISP can extract more of the value generated from OTTV.

Second, the results in Table 2 columns (5) - (8) suggest that steering consumers towards TV service is difficult. We find that, for the UBP policy we observe, the internet usage price increase was insufficient to steer many consumers from internet-only subscriptions into the bundle. This is true for both light and heavy users, despite

being sufficient to induce upgrades in plans. The relatively muted movement in either direction for TV subscriptions reveals the potential difficulty in influencing consumers along this margin and suggests that the ISP would need larger bundle discounts to steer more consumers, which would reduce the profitability of the steering-focused strategy.

Third, as the results in Table 3 show, UBP was effective in reducing overall usage, especially for consumers who did not upgrade their plan. Furthermore, consumers who choose to upgrade their internet tier did not use more data. Both these empirical findings suggest that the ISP's introduction of tiers, holding internet content constant, may not lead to pressure for an expanded network and higher ISP investment costs.

Fourth, the point estimates in Tables 4 and 5 show that consumers focused their usage on OTTV after UBP is introduced. This suggests that households favor streaming video over other types of internet content and therefore reductions in OTTV quality or its foreclosure would not be viewed favorably by consumers.

The tier upgrades and reduced usage represent outcomes that effectively reallocate some OTTV surplus from households to the ISP. Households' subscription changes over internet tiers and TV service increased the ISP's revenue by about 0.32%. Overage charges, which were paid by 4.7% of treated households, generated \$1.54 per household per month, which implies an additional 1.0% increase in overall ISP revenue. Despite the fact that some UBP outcomes can have negative impacts on some consumers, ISPs' incentives to encourage new content and expand their networks may be stronger when ISPs capture a greater share of total surplus from OTTV. Both the theoretical results and empirical findings suggest that the ISP can use both bundle discounts and usage tiers simultaneously in response to households' demand for streaming video. If the ISP's pricing strategies are more restricted, in particular through restrictions on its ability to price internet usage, the firm may perceive greater benefits in steering consumers toward conventional TV and restricting access to OTTV.

In addition to these direct impacts on the ISP and its subscribers, the pricing strategy impacts OTTV firms. The results in Table 5 suggest that Netflix and linear OTTV

subscriptions increased when UBP was introduced. While Netflix offers a mix of original programming and older TV shows and movies, other OTTV services like Hulu and SlingTV offer content that is even more similar to conventional TV. This suggests that consumers receive value from a range of different OTTV services, and an ISP effort to steer consumption toward one type of video entertainment (i.e., conventional TV) could limit the total surplus created by ISP subscriptions. Households' persistence in using streaming video, despite the UBP's introduction, is consistent with the ISP having opportunities to capture surplus from consumers' usage of these services, and therefore encourage their growth and improvement. At the same time, positive effects of UBP on OTTV subscription demand suggest that the OTTV firms themselves will continue improvements in developing and delivering content. The ISP's incentive and ability to direct consumers across video entertainment sources may play a prominent role in future antitrust or regulatory debates, such as evaluating vertical mergers between content and distributor firms, or more horizontal mergers involving firms that compete in either content production, video distribution, or both.

5 Conclusions

We study the impact of increasingly popular online video services on the telecommunications industry. OTTV represents an opportunity for ISPs in that streaming video increases the demand for internet subscription services, but the new applications also present several important challenges. OTTV improvements might reduce demand for ISPs' television services, and delivering internet content may substantially increase ISPs' costs.

We provide a model that describes some of the central incentives behind ISPs' pricing policies. We show that indeed the ISP might have an incentive to steer consumers away from OTTV, but that those incentives decline as the ISP has richer pricing tools that allow it to share in the gains generated by OTTV. We then use data from the introduction of UBP to understand consumer behavior and its implications for

the steering incentives. Consistent with the theoretical findings, the empirical results suggest that the ISP has an incentive to help improve internet access.

Understanding the incentives to steer is relevant for antitrust policy related to the telecommunications industry. In particular, evaluation of mergers between content and distribution firms presents a number of challenges. First, market boundaries may be difficult for regulators and antitrust authorities to identify because little evidence exists on consumers' willingness to substitute across conventional TV, streaming video, and other non-video internet applications. Our data show that consumers are willing to substitute among online entertainment sources and with conventional TV. Specifically, we find that cord cutters increase their usage of most online applications after dropping an ISP's TV service, and these increases are roughly proportional to usage levels prior to the subscription change. Thus, telecommunications antitrust analysis might need to consider broad market definitions that encompass many forms of digital entertainment, as well as the central role of ISPs in shaping how content is distributed and surplus is allocated.¹⁸

Second, antitrust authorities need to assess how existing or new vertical relationships may affect an ISP's incentives to introduce restrictive cross-licensing agreements or use price instruments to favor its own content over competitors'. The impact of these strategies depends on consumers' sensitivity to steering strategies. An ISP that is vertically integrated with a content-producing firm may foreclose some content from availability to consumers via a competing ISP. A price-based steering strategy with similar effects is "zero rating," which favors certain content by not counting its usage against a monthly allowance. Our estimates show that even blunt mechanisms like usage-based pricing can have important allocative consequences among consumers and various firms. Firms may perceive even greater benefits of zero rating than UBP because a targeted mechanism is likely to be more profitable.

More broadly, our results are also relevant for the net neutrality debate, in which

¹⁸We are not offering any specific market definition, nor did we conduct any formal analysis to conclude what the market should be. Such an analysis should be done on a case-by-case basis.

empirical evidence is rare. Net neutrality’s repeal provides ISPs more latitude to discriminate across types of internet traffic. While we do not observe source-specific discrimination in our analysis, our results are informative about ISPs’ incentives to discriminate when they have the opportunity. For example, we find that usage-based pricing’s primary impact is in inducing consumers to upgrade their internet tiers and continue using their preferred online applications (e.g., Netflix). If an ISP can successfully use tier premiums to extract some of the rents associated with OTTV innovations, it may not seek more targeted mechanisms to foreclose or diminish the attractiveness of OTTV. These incentives may change, however, as firms diversify and vertically integrate into media production, so more research is needed in this area.

There are several issues our model and empirical results do not address, and we leave for future research. While our model provides a useful framework for formalizing the steering incentives of ISPs, a richer specification is required to quantify the welfare implications of steering. Similarly, the model makes simplifying assumptions on the interaction between firms, for example competitive OTTV supply. Given the increasingly complex relationships between content providers and ISPs, and the evolving regulatory and antitrust environment, modeling and evaluating these policy issues is a fruitful area for future research.

References

- Armstrong, Mark (2006). “Competition in Two-Sided Markets.” *RAND Journal of Economics*, 37(3): 668–691.
- Armstrong, Mark (2013). “A More General Theory of Commodity Bundling.” *Journal of Economic Theory*, 148(2): 448–472.
- Barwick, Panle Jia, Parag A. Pathak and Maisy Wong (2017). “Conflicts of Interest and Steering in Residential Brokerage.” *American Economic Journal: Applied Economics*, 9(3): 191–222.

- Bourreau, Marc, Frago Kourandi and Tommaso Valletti (2015). “Net Neutrality with Competing Internet Platforms.” *Journal of Industrial Economics*, 63(1): 30–73.
- Choi, Jay Pil, Doh-Shin Jeon and Byung-Cheol Kim (2015). “Network Neutrality, Business Models, and Internet Interconnection.” *American Economic Journal: Microeconomics*, 7(3): 104–141.
- Choi, Jay Pil and Byung-Cheol Kim (2010). “Net Neutrality and Investment Incentives.” *RAND Journal of Economics*, 41(3): 446–471.
- Chu, Chenghuan Sean, Phillip Leslie and Alan Sorensen (2011). “Bundle-Size Pricing as an Approximation to Mixed Bundling.” *American Economic Review*, 101(1): 263–303.
- Cisco (2018). “Cisco Visual Networking Index (VNI) Complete Forecast Update, 2017–2022.” URL https://www.cisco.com/c/dam/m/en_us/network-intelligence/service-provider/digital-transformation/knowledge-network-webinars/pdfs/1213-business-services-ckn.pdf, retrieved: April 9, 2020.
- comScore (2017). “State of OTT, 2017.” URL <https://www.comscore.com/Insights/Presentations-and-Whitepapers/2017/State-of-OTT>, retrieved April 6, 2020.
- comScore (2018). “State of OTT, 2018.” URL <https://www.comscore.com/Insights/Presentations-and-Whitepapers/2018/State-of-OTT>, retrieved April 9, 2020.
- Crawford, Greg, Robin Lee, Mike Whinston and Ali Yurukoglu (2017). “The Welfare Effects of Vertical Integration in Multichannel Television Markets.” *Working Paper*.
- Crawford, Gregory and Matthew Shum (2007). “Monopoly Quality Degradation and Regulation in Cable Television.” *Journal of Law and Economics*, 50(1): 181–219.

- Crawford, Gregory and Matthew Shum (2015). “The Welfare Effects of Monopoly Quality Choice: Evidence From Cable Television Markets.” *Working Paper*.
- Crawford, Gregory and Ali Yurukoglu (2012). “The Welfare Effects of Bundling in Multichannel Television Markets.” *American Economic Review*, 102(2): 643–685.
- Dutz, Mark, Jonathan Orszag and Robert Willig (2009). “The Substantial Consumer Benefits of Broadband Connectivity for US Households.” *Internet Intervention Alliance Working Paper*.
- Economides, Nicholas and Benjamin Hermalin (2012). “The Economics of Network Neutrality.” *The RAND Journal of Economics*, 43(4): 602–629.
- Economides, Nicholas and Benjamin Hermalin (2015). “The Strategic Use of Download Limits by a Monopoly Platform.” *The RAND Journal of Economics*, 46(2): 297–327.
- Economides, Nicholas and Joacim Tag (2012). “Network Neutrality on the Internet: A Two-Sided Market Analysis.” *Information Economics and Policy*, 24(2): 91–104.
- Economides, Nicholas and Joacim Tag (2016). “Internet Regulation, Two-Sided Pricing, and Sponsored Data.” *Working Paper*.
- Edell, Richard and Pravin Varaiya (2002). *Providing Internet Access: What We Learn from INDEX*, volume Broadband: Should We Regulate High-Speed Internet Access? Brookings Institution.
- eMarketer (2019). “US Subscription Video Landscape 2019.” URL <https://www.emarketer.com/content/us-subscription-video-landscape-2019>, retrieved April 9, 2020.
- Gans, Joshua (2015). “Weak Versus Strong Net Neutrality.” *Journal of Regulatory Economics*, 47(2): 183–200.
- Gentzkow, Matthew (2007). “Valuing New Goods in a Model with Complementarity: Online Newspapers.” *American Economic Review*, 97(3): 713–744.

- Goetz, Daniel (2016). “Competition and Dynamic Bargaining in the Broadband Industry.” *Princeton University Working Paper*.
- Goolsbee, Austan and Peter Klenow (2006). “Valuing Products by the Time Spent Using Them: An Application to the Internet.” *American Economic Review P&P*, 96(2): 108–113.
- Greenstein, Shane and Ryan McDevitt (2011). “The Broadband Bonus: Estimating Broadband Internet’s Economic Value.” *Telecommunications Policy*, 35(7): 617–632.
- Greenstein, Shane, Martin Peitz and Tommaso Valletti (2016). “Net Neutrality: A Fast Lane to Understanding the Tradeoffs.” *Journal of Economic Perspectives*, 30(2): 127–150.
- Grubb, Michael (2015). “Consumer Inattention and Bill-Shock Regulation.” *Review of Economic Studies*, 82(1): 219–257.
- Grubb, Michael and Matthew Osborne (2015). “Cellular Service Demand: Biased Beliefs, Learning, and Bill Shock.” *American Economic Review*, 105(1): 234–271.
- Hagemann, Garrett (2017). “Upstream Quantity Discounts and Double Marginalization in the New York Liquor Market.” *Working Paper*.
- Hitte, Loran and Prasanna Tambe (2007). “Broadband Adoption and Content Consumption.” *Information Economics and Policy*, 74(6): 1637–1673.
- Ho, Kate and Robin Lee (2019). “Equilibrium Insurer-Provider Networks: Bargaining and Exclusion in Health Care Markets.” *American Economic Review*, 109(2).
- Lambrecht, Anja, Katja Seim and Bernd Skiera (2007). “Does Uncertainty Matter? Consumer Behavior Under Three-Part Tariffs.” *Marketing Science*, 26(5): 698–710.
- Lee, Robin and Tim Wu (2009). “Subsidizing Creativity through Network Design: Zero-Pricing and Net Neutrality.” *Journal of Economics Perspectives*, 23(3): 61–76.

- Liebman, Eli (2017). “Bargaining in Markets with Exclusion: An Analysis of Health Insurance Networks.” *Working Paper*.
- Malone, Jacob, Aviv Nevo and Jonathan Williams (2016). “A Snapshot of the Current State of Residential Broadband Networks.” *NET Institute Working Paper No. 15-06*.
- Malone, Jacob, John Turner and Jonathan Williams (2014). “Do Three-Part Tariffs Improve Efficiency in Residential Broadband Networks?” *Telecommunications Policy*, 38(11): 1035–1045.
- MarketWatch (2018). “Over 5 million U.S. consumers will cut the cord in 2018, survey says.” URL <https://www.marketwatch.com/story/over-5-million-us-consumers-will-cut-the-cord-in-2018-survey-says-2018-07-17>, retrieved April 6, 2020.
- McManus, Brian (2007). “Nonlinear Pricing in an Oligopoly Market: the Case of Specialty Coffee.” *RAND Journal of Economics*, 38: 512–532.
- Miravete, Eugenio (2003). “Choosing the Wrong Calling Plan? Ignorance and Learning.” *American Economic Review*, 93(1): 297–310.
- Mussa, Michael and Sherwin Rosen (1978). “Monopoly and Product Quality.” *Journal of Economic Theory*, 18(2): 301–317.
- Nevo, Aviv, John Turner and Jonathan Williams (2016). “Usage-Based Pricing and Demand for Residential Broadband.” *Econometrica*, 84(2): 411–443.
- Nielsen (2015). “The Total Audience Report: Q4 2014.” URL <https://www.nielsen.com/us/en/insights/report/2015/the-total-audience-report-q4-2014>, retrieved April 3, 2020.
- Prince, Jeffrey and Shane Greenstein (2017). “Measuring Consumer Preferences for Video Content Provision via Cord-Cutting Behavior.” *Journal of Economics & Management Strategy*, 26(2): 293–317.

- Raval, Devesh and Ted Rosenbaum (2017). “Payer Steering and Patient Welfare: Evidence from Medicaid.” *Working Paper*.
- Reggiani, Carlo and Tommaso Valletti (2016). “Net Neutrality and Innovation at the Core and at the Edge.” *International Journal of Industrial Organization*, 45(1): 16–27.
- Rosston, Gregory, Scott Savage and Bradley Wimmer (2013). “Effect of Network Unbundling on Retail Price: Evidence from the Telecommunications Act of 1996.” *Journal of Law and Economics*, 56(2): 487–519.
- Sidak, Gregory (2006). “A Consumer-Welfare Approach to Network Neutrality Regulation of the Internet.” *Journal of Competition Law & Economics*, 2(3): 349–474.
- Tudon, Jose (2018). “Congestion v. Content Provision in Net Neutrality: The Case of Amazon’s Twitch.tv.” *Working Paper*.
- Varian, Hal (2002). *The Demand for Bandwidth: Evidence from the INDEX Experiment*, volume Broadband: Should We Regulate High-Speed Internet Access? Brookings Institution.
- Wu, Tim (2003). “Network Neutrality, Broadband Discrimination.” *Journal of Telecommunications and High Technology Law*, 1(2): 141–178.
- Zenith (2019). “Online video viewing to reach 100 minutes a day in 2021.” URL <https://www.zenithmedia.com/online-video-viewing-to-reach-100-minutes-a-day-in-2021>, retrieved April 6, 2020.

Appendix A

A.1 Internet Usage Changes when TV Service is Added

To complement the analysis of Section 2.3, we examine the behavior of households that added a conventional TV subscription. In Figure A1 panel (a), we track the internet usage changes for households who added video to their subscriptions. Internet usage falls significantly after this subscription change. The reduction’s magnitude (12%) is about half the magnitude of the increase in internet usage that occurs after a household cuts the cord.

[FIGURE A1 HERE]

In panel (b) of Figure A1, we show how households’ usage of major internet services changes after adding a TV subscription. Average usage of some types of applications like Browsing, Music, and Gaming increase slightly during our sample, while Video streaming falls significantly. The largest absolute and percentage reduction in usage is in Netflix usage.

A.2 Additional Model Discussion

In this section, we augment the discussion in the main text, and provide a more detailed description of how consumers’ choices would be affected when the ISP introduces usage tiers. In practice, the tiers may vary in terms of usage allowances, overage prices, or connectivity speeds. We focus on a case in which the ISP, facing $\delta > 0$, shifts from having no tiers and no usage allowances to having both. This is similar to the situation in our data, and it also resembles the implementation of the tiering and UBP strategy implemented by numerous North American ISPs.

In Figure A2 we demonstrate the effect of the ISP introducing a “low-usage” tier which is subject to the usage allowance κ while a “high-usage” tier has no usage limit. The same tiers are available to both internet-only and bundle subscribers (i_L and b_L for low-usage tiers, and i_H and b_H for high usage). We impose a set of prices that

facilitate reading the different regions of Figure A2; the optimal prices and cap value would change the sizes and perhaps shapes of some regions.

Panel (a) of Figure A2 provides an initial distribution of consumers across subscription options, before any tiers or caps are offered. As in Figure 3 Panel (a), with $\delta > 0$ a significant share of consumers select the i subscription and satisfy their video entertainment tastes with OTTV. The introduction of usage tiers splits i and b subscribers into a new collection of actions, illustrated in Figure A2 Panel (b).

[FIGURE A2 HERE]

Former subscribers to the initial unlimited internet service (i) may update their subscription and usage choices in several ways. Some consumers, in Figure A2 Panel (b)'s area I, will accept the usage cap κ and remain internet-only. Consumers with a stronger taste for internet usage, whether for video- or non-video entertainment, may “upgrade” their internet subscription to i_H ; these consumers are in Panel (b)'s area II. From the ISP's perspective, the tier upgrade is a way to have area II's consumers pay a greater price for internet service that is equal to what the consumers received previously. Consumers with relatively strong values of v_2 switch from i into the bundle (areas III and IV). Of these consumers, those with high values of v_1 pay for a tier upgrade in addition to TV service (area IV). When consumers switch to the bundle, they receive video entertainment through TV, so their OTTV usage falls to zero. Finally among initial i subscribers, some will cancel their subscriptions completely (area V) because their relatively strong taste for video entertainment and weak taste for non-video internet means that capped internet, at the present price, is worth less than the outside option. In addition to these margins for former i subscribers, some bundle subscribers with strong internet tastes (in area VI) will opt for b_H so that they can consume internet without a usage limit.

A.3 Evaluation of Pre-Treatment Trends

Our empirical analysis compares outcomes in the treated and control markets. In this section, we investigate whether the markets' trends prior to UBP were meaningfully different from each other. We examine trends in a collection of variables that are closely related to our main empirical results.

For four variables related to ISP subscriptions, we ask whether there are differences in subscription transitions during the pre-policy period between the treated and control markets.¹⁹ We use one observation per household, and we estimate a Probit model in which the dependent variable is the household's realized decision regarding: upgrading internet tier, downgrading internet tier, adding a TV subscription, or dropping a TV subscription. For each dependent variable, the sole explanatory variable is *PreTreatment*, an indicator for whether the household is in the treated market. If treated households are more likely to, for example, add TV, then this tendency will be apparent in the treated market before UBP is announced. In the first four rows of Table A1, we report estimation results for the coefficient on *PreTreatment*. We find that all four models show very small or zero difference between the treatment and control markets. The coefficients on *PreTreatment* are significantly different from zero in three models, but the coefficient magnitudes are not economically significant. The pre-policy period is about one quarter of the duration of the announcement and treatment periods, so the coefficients in Table A1 should be multiplied by four for comparison to the results presented in Section 4.2. Even after this adjustment, the coefficients on Table A1 imply very small differences between the treatment and control markets. For example, the tier upgrading rate in the treated market is about an order of magnitude larger after UBP was announced relative to the pre-policy period.

For six additional variables, we examine weekly data on households' use of the internet and applications on it. The variables we examine are: the natural log of total internet usage; the shares of all internet usage devoted to video, devoted to Netflix,

¹⁹Subscription choices are made relatively infrequently, so we cannot investigate a trend within the pre-policy period.

and devoted to linear OTTV; and indicators for whether the household used 0.1GB or more of Netflix and (separately) linear OTTV on any day of the observed week. For each variable, we sum the households' activity to the market level, and then we regress the market-level measure on a market-level dummy variables, a linear time trend, and the time trend interacted with *PreTreatment*. In Table A1 we report the coefficient estimates for the interaction of the time trend and *PreTreatment*. In all cases but one (share of internet usage for linear OTTV), the coefficients are not significantly different from zero. In the case of linear OTTV usage, the coefficient estimate is essentially zero and therefore economically insignificant.

[TABLE A1 HERE]

6 Tables and Figures

Table 1: Adoption and Usage of OTTV Applications

	Ever Active Users (%)	Daily Usage (GB)	
		All HHs	Active Users
Internet-only HHs			
HBO Go	14.1	0.03	0.53
Hulu	23.2	0.13	1.07
Netflix	77.6	1.75	2.86
SlingTV	3.3	0.05	3.00
YouTube	88.7	1.06	1.49
Bundle HHs			
HBO Go	11.8	0.02	0.49
Hulu	13.2	0.04	0.85
Netflix	67.9	1.10	2.27
SlingTV	0.83	0.01	2.81
YouTube	74.3	0.69	1.26

Notes: A household is an active user of an application if it uses 3GB or more during a month. The “Ever Active User” column provides the share of households that are identified as active users during one or more months of the sample period. The daily usage statistics for active users are calculated for months when usage is 3GB or more. The top panel reports these statistics for internet-only households, and the bottom panel contains households with both internet and TV subscriptions.

Table 2: The Effect of UBP on Take-up of ISP Subscriptions

	Upgrade Tier		Downgrade Tier		Add TV		Drop TV	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treatment	0.030*** (0.002)		0.002 (0.001)		0.003 (0.004)		0.001 (0.002)	
Treatment × Share0-50		0.016*** (0.002)		0.002* (0.001)		0.006* (0.003)		0.000 (0.002)
Share50-100		0.053*** (0.003)		-0.001 (0.001)		-0.005 (0.005)		0.008*** (0.001)
Share100+		0.112*** (0.011)		-0.006*** (0.002)		0.004 (0.006)		-0.006** (0.003)
Observations	409,073	409,073	383,959	383,959	114,435	114,435	304,742	304,742

Notes: The table reports marginal effects from Probit regressions. The dependent variable, listed in the first row of the table, is an indicator for whether the household makes the specified subscription transition between the announcement of UBP and the end of the sample. An observation is a household and the sample includes all households eligible to make the specified subscription change (e.g, in columns (1) and (2) all households are included except those already at the highest tier). $Share_N$ is a dummy equal to one if in the pre-policy period the household used $N\%$ of the their internet tier's treatment-period allowance. All regressions include as controls the $Share_h$ variables, position in the monthly total internet usage distribution (discretized into deciles), pre-policy growth in total usage, usage deciles interacted with pre-policy growth, OTTV subscription indicators, internet tier and TV subscription dummies. Standard errors, clustered by market, are in parentheses, with stars indicating the following significance levels: * ($p < 0.10$), ** ($p < 0.05$), *** ($p < 0.01$).

Table 3: The Effect of UBP on Internet Usage

	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	-0.110*		-0.121*		-0.044	
	(0.045)		(0.046)		(0.036)	
Treatment ×						
Share0-50		-0.107*		-0.111*		-0.085
		(0.048)		(0.048)		(0.047)
Share50-100		-0.104**		-0.154**		0.104**
		(0.035)		(0.036)		(0.037)
Share100+		-0.190***		-0.329***		-0.081*
		(0.035)		(0.035)		(0.037)
Constant	3.441***	3.441***	3.432***	3.432***	3.647***	3.647***
	(0.020)	(0.020)	(0.021)	(0.021)	(0.024)	(0.023)
Sample	All HHs	All HHs	Upgr=N	Upgr=N	Upgr=Y	Upgr=Y
Observations	3,789,621	3,789,621	3,625,398	3,625,398	164,223	164,223

Notes: The table reports estimates from linear regressions. An observation is a household-month. The dependent variable is log internet usage in gigabytes. *ShareN* is a dummy equal if in the pre-policy period the household used N% of the their internet tier's treatment-period allowance. Additional controls include: household fixed effects; a heterogeneous time trend which interacts deciles of total usage distribution with a cubic time trend; and the *Announcement_{ht}* variable, which is interacted with *ShareN* in the even-numbered columns. Standard errors, clustered at the market level, are reported in parentheses, with stars indicating the following significance levels: * ($p < 0.10$), ** ($p < 0.05$), *** ($p < 0.01$).

Table 4: The Effect of UBP on Online Video Usage

	Video Share			Netflix Share		LinOTTV Share	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treatment	0.012 (0.020)						
Treatment \times Share0-50		0.009 (0.016)	0.028 (0.025)	0.018 (0.014)	0.039 (0.026)	0.001** (0.000)	0.003* (0.001)
Share50-100		0.031 (0.045)	0.018 (0.034)	0.025 (0.039)	0.027 (0.034)	0.003 (0.002)	0.007 (0.005)
Share100+		0.020 (0.048)	0.021 (0.036)	0.018 (0.040)	0.024 (0.034)	0.004 (0.002)	0.021* (0.006)
Constant	0.383*** (0.003)	0.383*** (0.003)	0.630*** (0.003)	0.158*** (0.005)	0.394*** (0.006)	0.008*** (0.000)	0.161*** (0.002)
Observations	2,508,084	2,508,084	1,006,732	2,508,084	986,121	2,508,084	146,345

Notes: The table reports estimates from linear regressions. An observation is a household-month. The dependent variable, depending on the column is the share of video/Netflix/Linear OTTV out of total internet usage. The sample in columns (1), (2), (4) and (6), includes all households-months, while other columns the sample includes household-months with a Netflix subscription (column (5)), linear OTTV subscription (column (7)), or both (column (3)). $ShareN$ is a dummy equal if in the pre-policy period the household used N% of the their internet tier's treatment-period allowance. Additional controls include: household fixed effects; a heterogeneous time trend which interacts deciles of total usage distribution with a cubic time trend; and the $Announcement_{ht}$ variable, which is interacted with $ShareN$ in the even-numbered columns. Standard errors, clustered at the market level, are in parentheses, with stars indicating the following significance levels: * ($p < 0.10$), ** ($p < 0.05$), *** ($p < 0.01$).

Table 5: The Effect of the Price Change on Take-up of OTTV Subscriptions

	Netflix			LinOTTV		
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	0.033** (0.014)			0.002 (0.007)		
Treatment × Share0-50		0.030 (0.020)	0.008** (0.004)		0.002 (0.007)	-0.015 (0.014)
Share50-100		-0.047 (0.054)	0.016*** (0.001)		-0.001 (0.010)	0.028** (0.013)
Share100+		-0.093 (0.064)	-0.023*** (0.002)		-0.003 (0.012)	-0.047*** (0.006)
Sample	All	New	Keep	All	New	Keep
Observations	278,676	153,679	124,997	278,676	260,247	18,429


Notes: Each column contains marginal effects from a probit regression. Each observation is a single household. The dependent variable is an indicator for whether the household subscribes to the OTTV service specified in the first row of the table during the treatment period. All regressions include as controls the uninteracted $Share_h$ variables, position in the monthly total internet usage distribution (discretized into deciles), pre-policy growth in total usage, usage deciles interacted with pre-policy growth, OTTV subscription indicators, internet tier and TV subscription dummies. Standard errors are in parentheses, with stars indicating the following significance levels: * ($p < 0.10$), ** ($p < 0.05$), *** ($p < 0.01$).

Table A1: Pre-policy Regressions

	Coefficient (SE)	p-value
Upgrade Tier	0.001*** (0.000)	0.000
Downgrade Tier	-0.001* (0.001)	0.098
Add TV	0.001* (0.000)	0.054
Drop TV	-0.000 (0.000)	0.679
Log Total GB	-0.010 (0.009)	0.253
Video Share	0.001 (0.001)	0.460
Netflix Share	0.000 (0.001)	0.995
LinOTTV Share	0.000** (0.000)	0.039
Netflix Subscriptions	-0.000 (0.002)	0.896
LinOTTV Subscriptions	0.000 (0.000)	0.657

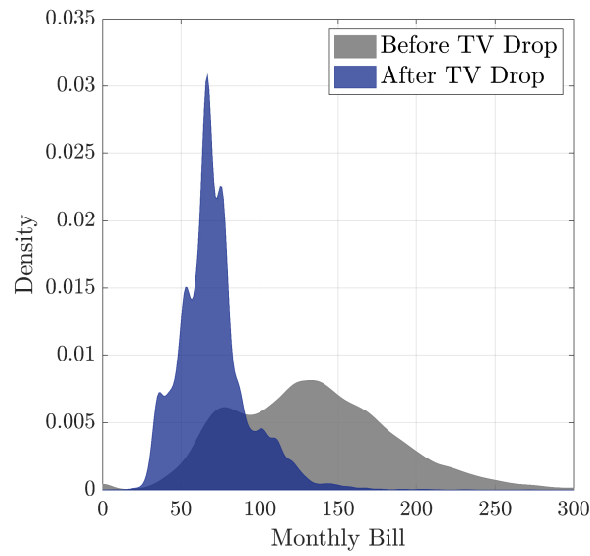
Notes: Each row is a regression with dependent variable indicated in the first column. The first four rows are subscriber-level regressions of the binary outcomes from Table 2 on a treatment indicator. The final six rows are regressions of the continuous outcome variables from Tables 3-5 on a time trend, market dummies, and the interaction of the time trend and treatment indicator at the market-week level. The first eight variables are exactly as defined in Section 4 regressions. The Netflix and LinOTTV subscriptions variables are defined as the share of households that are active users of the indicated application using the same threshold as Table 5 (usage exceeding 0.1 GB/day), except at the weekly level. The second column reports the coefficient on the treatment indicator (first four rows) and treatment-time trend interaction (final six rows), with standard errors in parentheses. p values are reported in the third column.

Figure 1: Cord-cutter Usage Changes



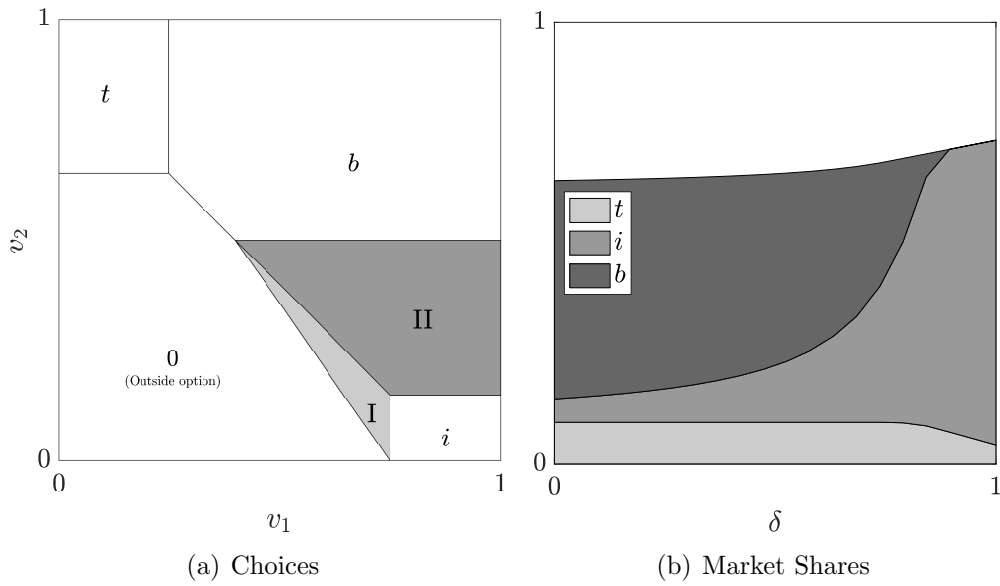
Notes: Panel (a) summarizes the average daily usage (GBs) of internet subscribers in the eight weeks before and after dropping TV services (upper plots) in comparison to all other internet subscribers (lower plots). The dots are average daily usage, the dashed line is average usage across before and after the switching event (with two weeks on either side omitted), and the solid line is a global local-linear regression fit of the data. The shaded areas are 95% confidence intervals. Panel (b) summarizes the average daily usage (GBs) of internet traffic types by cord cutters in the eight weeks before and after dropping TV services (two weeks on either side of the switching event are omitted). For each traffic type, the percent change in usage level between the two periods is shown above the two bars.

Figure 2: Estimated Total Subscriber Monthly Payments with OTTV



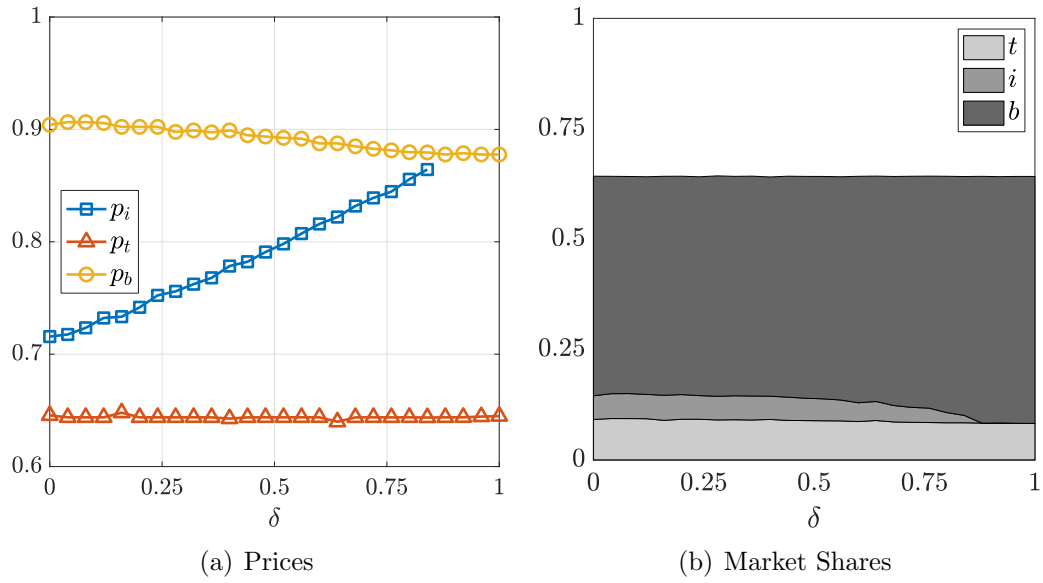
Notes: This figure shows the distribution of estimated monthly payments by cord cutters in the month before and after dropping pay TV. The total monthly operator bill is observed in the ISP billing data. Monthly OTTV expenditures are estimated using 2015 subscription fees for each of the video services observed in the usage data.

Figure 3: Consumer Response to Changes in δ



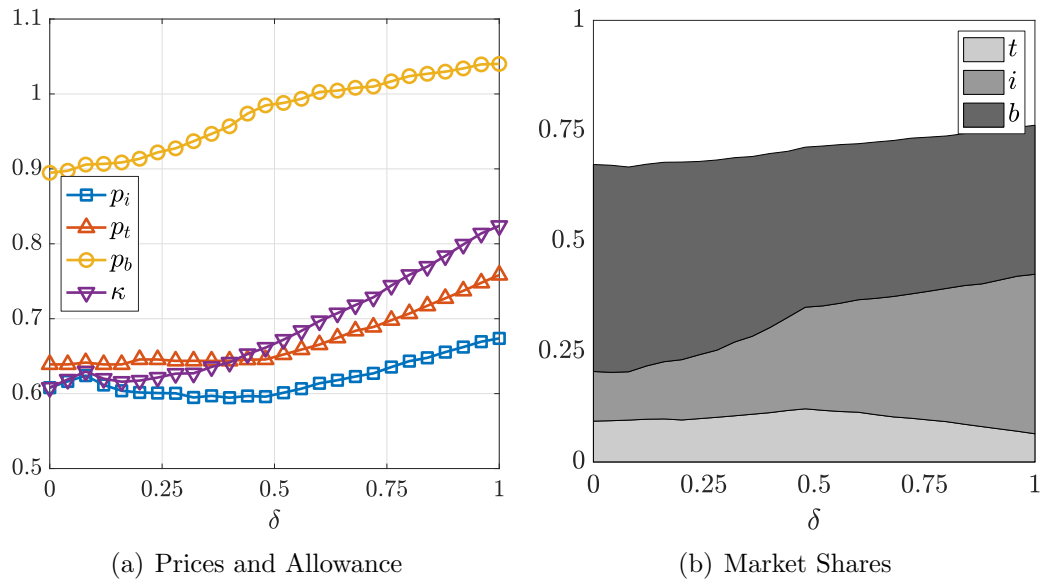
Notes: These figures illustrate consumer choices for different levels of δ , holding prices fixed. In panel (a) we present consumer choices for δ equal to 0 and 0.7 holding prices fixed at $(p_i, p_t, p_b) = (0.75, 0.65, 0.9)$, which are the profit-maximizing prices for $\delta = 0$, $\beta = 1.25$, and $\gamma = 0.2$. In panel (b), we present market shares when δ increases from 0 to 1 and prices are fixed at $(p_i, p_t, p_b) = (0.75, 0.65, 0.9)$.

Figure 4: Optimal Prices for Varying Values of δ



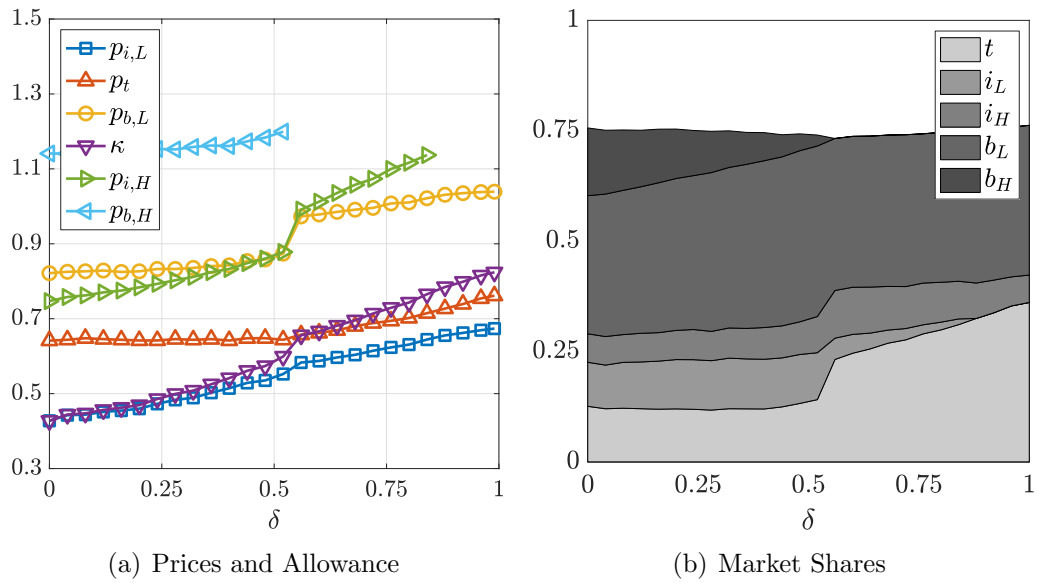
Notes: Panel (a) displays the ISP's profit-maximizing prices as δ varies from 0 to 1. Panel (b) provides the market shares resulting from these prices. We compute prices and shares numerically using the assumptions described in the text.

Figure 5: Optimal Prices and Allowance with Usage Limits for Varying Values of δ




Notes: Panel (a) displays the ISP's profit-maximizing prices and allowance when the ISP sets a usage allowance, as δ varies from 0 to 1. Panel (b) summarizes the change in market shares at each optimal price combination.

Figure 6: Optimal Tiered Prices and Allowance for Varying Values of δ



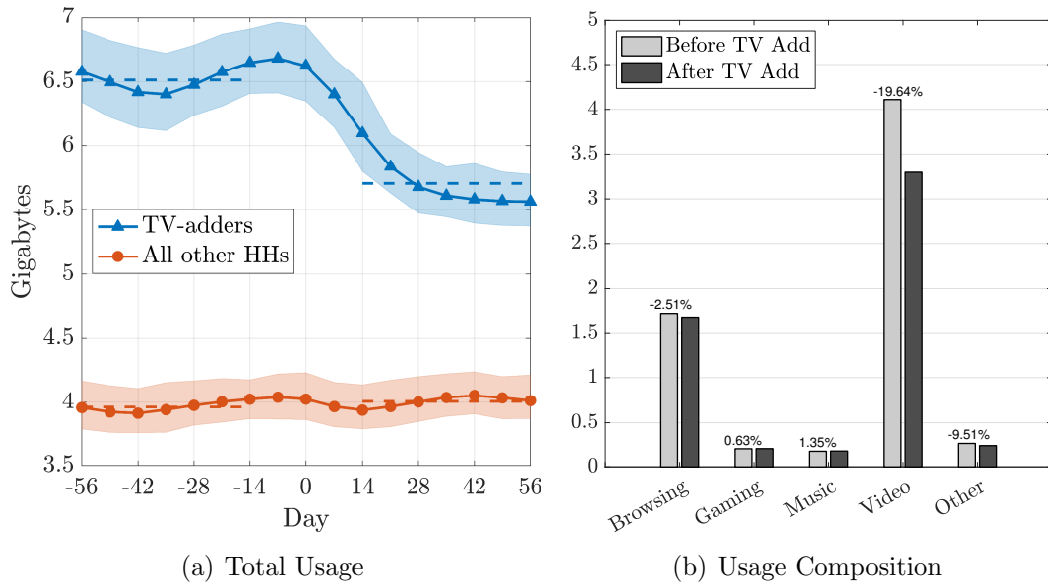
Notes: Panel (a) depicts the ISP's optimal tier prices and usage allowance as δ varies from 0 to 1. Panel (b) summarizes the change in market shares at each optimal price and allowance combination.

Figure 7: Profits by Pricing Strategy for Varying Values of δ



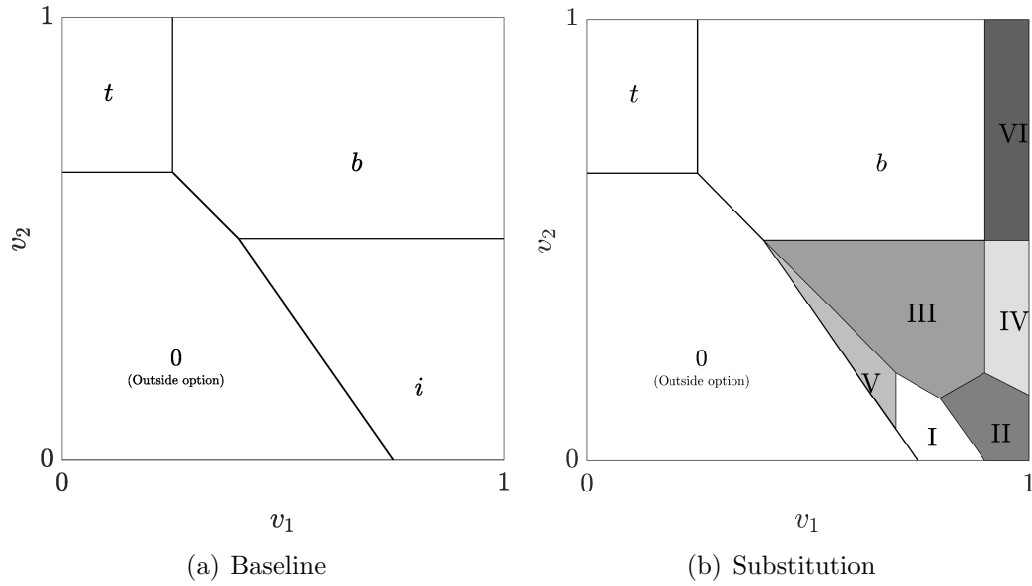
Notes: This figure shows ISP profits when it uses optimal bundle prices, optimal prices and usage cap, and optimal prices and cap with usage tiers.

Figure A1: Usage Changes when TV Service is Added



Notes: Panel (a) summarizes the average daily usage (GBs) of internet subscribers in the eight weeks before and after adding TV services (upper plots) in comparison to all other internet subscribers (lower plots). The dots are average daily usage, the dashed line is average usage across before and after the switching event (with two weeks on either side omitted), and the solid line is a global local-linear regression fit of the data. The shaded areas are 95% confidence intervals. Panel (b) summarizes the average daily usage (GBs) of online traffic types by TV-adders in the eight weeks before and after adding TV services (2 weeks on either side of the switching event omitted). For each traffic type, the percent change in usage level between the two periods is shown above the two bars.

Figure A2: Effect of Tiers and Allowances



Notes: This figure shows the effect of the introduction of a tier with a usage allowance on subscription choices. Throughout, δ is fixed at 0.7. Market shares are first plotted for prices $(p_i, p_t, p_b) = (0.75, 0.65, 0.9)$. Next, a usage allowance is placed on the original internet tier, and a new premium internet tier is introduced with no allowance. The new prices are $(p_{i,L}, p_{i,H}, p_t, p_{b,L}, p_{b,H}) = (0.75, 0.85, 0.65, 0.9, 1.0)$. Each shaded region depicts a set of consumer types that makes a particular subscription change. The usage allowances κ is equal to 0.8.